LBNL-182764

ERNEST ORLANDO LAWRENCE BERKELEY NATIONAL LABORATORY

Loss Aversion and Time-Differentiated Electricity Pricing

C. Anna Spurlock

Energy Analysis and Environmental Impacts Division
Lawrence Berkeley National Laboratory

June 2015

Support for this work was provided to Lawrence Berkeley National Laboratory by the U.S. Department of Energy under Contract No. DE-AC02-05CH11231.

Loss Aversion and Time-Differentiated Electricity Pricing

C. Anna Spurlock*

June 9, 2015

Abstract

I develop a model of loss aversion over electricity expenditure, from which I derive testable predictions for household electricity consumption while on combination time-of-use (TOU) and critical peak pricing (CPP) plans. Testing these predictions results in evidence consistent with loss aversion: (1) spillover effects - positive expenditure shocks resulted in significantly more peak consumption reduction for several weeks thereafter; and (2) clustering - disproportionate probability of consuming such that expenditure would be equal between the TOU-CPP or standard flat-rate pricing structures. This behavior is inconsistent with a purely neoclassical utility model, and has important implications for application of time-differentiated electricity pricing.

*Lawrence Berkeley National Laboratory, 1 Cyclotron Road, Mail Stop 90R4000 Berkeley, CA 94720 USA, caspurlock@lbl.gov. A previous version of this research appeared in my doctoral dissertation from University of California, Berkeley entitled "Essays in Energy Economics." I would like to thank Peter Berck, Meredith Fowlie, Stefano DellaVigna, Sofia Berto Villas-Boas, Catherine Wolfram, and Koichiro Ito for their invaluable advice and comments. I am very grateful to the California Energy Commission, and particularly David Hungerford, for giving me permission to use the data from the California Statewide Pricing Pilot, and to the Energy Institute at Haas for facilitating my access to these data. All mistakes are my own. Support for this work was provided to Lawrence Berkeley National Laboratory by the U.S. Department of Energy under Contract No. DE-AC02-05CH11231.

Charging a static price for retail electricity in the face of wholesale price volatility and demand fluctuations can result in electricity shortages, as well as over-investment in, and under-utilization of, production capacity in the long run (Borenstein, Jaske and Rosenfeld 2002; Joskow and Wolfram 2012). Time-differentiated pricing mechanisms strengthen the connection between wholesale and retail prices, particularly during high demand hours. The simplest is time-of-use (TOU) pricing, wherein a low price is charged during off-peak hours, and a higher price charged during peak hours. A further extension of this concept is critical peak pricing (CPP), wherein the utility has the ability to charge a very high price for peak consumption during a limited number of critical days. This improves upon TOU pricing by providing the utilities with a way to respond when demand projections approach the capacity constraint of the system.[1]

While there are relatively few residential electricity customers currently enrolled in time-based pricing (approximately 1 percent as of 2010), the explosion of smart-meter technology – facilitating the wide scale implementation of such pricing options – in recent years has increased drastically (Joskow and Wolfram 2012). In addition, there has be significant recent policy shifts paving the way toward wide-scale implementation of time varying pricing. In particular, in California Senate Bill 695 has ruled that the California Public Utilities Commission (CPUC) can transition residential customers onto time varying rates starting in 2013. To this end, the CPUC has initiated a rule making (R.12-06-013) designed to explore the way in which this transition can take

[1] Other examples of time-based pricing structures include real time pricing (RTP), wherein the electricity price varies continuously throughout the day in response to wholesale price fluctuations; and peak time rebates (PTR), which are similar to a CPP tariff except the incentive to reduce peak consumption during critical days comes in the form of rebates for forgone consumption rather than higher prices.

place.

Understanding the underlying patterns of consumer behavior when placed on these rates is increasingly important both to inform policy makers in understanding customer impacts when faced with time-based pricing, and to inform electricity utilities in the design and planning of such rate plans.

An important aim of time-differentiated pricing is to reduce peak demand on the handful of days in which the wholesale price for electricity is are extremely high (Joskow and Wolfram 2012), and CPP mechanisms tend to be the most effective at achieving this goal (Faruqui 2010), particularly relative to other time-differentiated pricing mechanisms such as TOU and peak time rebate/critical peak rebate (PTR/CPR). While time-differentiated tariffs influence consumption through a simple price response, the psychology and economics literature – and particularly the concept of loss aversion – may contribute insights into why CPP specifically is so effective.[2] Loss aversion, a feature of reference-dependent utility, posits that consumers experience a larger impact to their utility from a loss relative to a gain. Loss aversion is relevant to time-based pricing because, as prices change over time, consumers incur expenditures higher than they are used to (a loss) in some bill periods, and lower than they are used to (a gain) in others. Loss aversion predicts that consumers will modify their consumption in predictable and policy-relevant ways in order to avoid high losses and to enjoy gains.

In this paper I outline a model of loss aversion over electricity expenditure, test predictions from this model, and find evidence consistent with loss

[2]There are other possible explanations for this that cannot be explicitly addressed with the data used in this paper. For example, price awareness or salience may be a factor, at least in the relative effectiveness of CPP to TOU, if not to PTR. Jessoe and Rapson (2014) discuss the role of awareness and salience in the context of TOU rates with and without in-home information displays.

aversion in the electricity consumption behavior of households participating in a TOU-CPP pricing experiment. One prediction is that conditional on electricity prices, households will reduce consumption of high-priced electricity measurably more if they are more likely to be in the loss domain of their value function rather than in the gain domain, and indeed I find evidence that consumers reduce their consumption in high-priced peak hours disproportionately more if they have experienced positive expenditure shocks already within a given bill period. A second prediction is that consumption patterns will disproportionately cluster bill-period expenditure at the kink in the value function where there is zero loss or gain. Again, I show evidence of disproportionate clustering at the kink in the reference-dependent value function (where expenditure is equal to what it would have been on standard prices), particularly when prices are structured in such a way as to place households close to the kink, and to skew their expenditure into the loss domain.[3]

Previous empirical work has demonstrated evidence of loss aversion in various facets of economic behavior. Two excellent review articles discussing applications of loss aversion are Camerer (2000) and DellaVigna (2009). The role of loss aversion has been explored in several settings, including the stock market, explanations for the equity premium puzzle (Benartzi and Thaler 1995; Barberis, Huang and Santos 2001) and the disposition effect (Odean 1998; Barberis and Xiong 2009; Li and Yang 2013); labor supply (Camerer et al. 1997; Farber 2005, 2008; Fehr and Goette 2007); employment (Mas 2006);

[3]In this work I examine one of many possible intersections between the literature on demand-side management of electricity markets on the one hand, and the psychology and economics literature on the other. Some previous empirical studies have also bridged these two fields. A few examples are: Hartman, Doane and Woo (1991) who demonstrate evidence of a wedge between consumer willingness to accept and willingness to pay for changes in electricity service reliability; and Ayres, Raseman and Shih (2009), Allcott (2011), Costa and Kahn (2013) and Allcott and Rogers (2014) who study social norms and electricity conservation using the Opower billing mechanism.

consumer choice, as an explanation for status quo bias in health insurance choice (e.g., Samuelson and Zeckhauser 1988; Sydnor 2010) or the choice to sell at a loss or not in the housing market (Genesove and Mayer 2001); and finally, consumption behavior, as an explanation for asymmetric price elasticities (Hardie, Johnson and Fader 1993). This paper contributes to this literature by demonstrating evidence consistent with loss aversion in a new and previously unexplored setting, namely household electricity consumption behavior and time-differentiated pricing. The results in this paper are most related to previous work on asymmetric price elasticities, however Hardie, Johnson and Fader (1993) focus on a discrete brand choice framework, while this paper explores implications for the sensitivity of quantity to price, based on the probability of experiencing a loss in expenditure, on a single good from one supplier.

This paper proceeds as follows. Section I presents the model and derives the testable predictions; section II discusses the data; section III presents the estimation strategy and results; section IV discusses alternative hypotheses; and section V concludes.

I Model

I develop a model of utility and demand with reference-dependent preferences over expenditure on electricity. I use an extension of the original Kahneman and Tversky (1979) Prospect Theory model developed by Sugden (2003) and Köszegi and Rabin (2006), in which utility is derived not only from outcomes relative to a reference point, but from the level of the outcome as well.[4] I

[4]While Kahneman and Tversky's original model consists of four features (reference-dependence; loss aversion; risk aversion over gains and risk seeking over losses; and differential probability weighting), I follow the example of much of the empirical literature in

assume that consumers "narrowly bracket" their sense of gains and losses at the bill period level.[5] This means that consumers experience either a gain or loss over the electricity expenditure incurred during their current bill period in addition to the direct utility they obtain from the consumption of electricity.[6]

The nature and malleability of reference points is a substantial area of research, with many approaches and context-specific hypotheses – a few examples of work exploring this issue are Köszegi and Rabin (2006), Arkes et al. (2008), Baucells, Weber and Welfens (2011), and De Giorgi and Post (2011). I assume the reference point in this case is based off of the standard time-invariant electricity prices that the consumers used to face prior to the CPP pilot. In particular, I assume that the electricity expenditure reference point for a given consumer on a TOU-CPP pricing tariff for a given bill period is the amount she would have paid for her realized level of consumption if she had been charged the standard time-invariant price for electricity instead. The intuition is that households have a sense of how much it usually costs them to use a set of electricity services, and will base their reference point on how much they would usually expect to pay during the bill cycle for their chosen set of household behaviors and activities.[7]

this area and incorporate only reference-dependence and loss aversion (DellaVigna 2009).

[5] I do a robustness check of the primary results from table 4 wherein I artificially shift the definition of the bill period forward by 14 days. These results are presented in online Appendix G. I do this in order to test to what extent households are actually narrowly bracketing within the bill period, and find that this assumption is too restrictive given the results in this robustness check. However, for the purposes of the exposition of the model, I maintain this assumption, but simply caveat that consumers are possibly using heuristics to approximate this, given that there is a good chance they do not know exactly when their bill periods start and stop.

[6] Note that I use the terms *consumer* and *household* interchangeably.

[7] I make the assumption that consumers have a reference point over bill period electricity expenditure (instead of daily expenditure, for example) because the most salient expenditure feedback most consumers receive regarding their electricity expenditure is their bill. It's not clear that consumers would necessarily be aware of the specific price differential on a daily or hourly level.

The following variables are taken as given: temperature during peak and off-peak hours along with other determinants of demand captured in the vector $\boldsymbol{x}_{it}$; the vector of off-peak and peak current electricity prices $\boldsymbol{p}_{it} = (p_{op,it}, p_{p,it})'$; the time-invariant "reference" prices $\boldsymbol{p}_{r,it} = (p_{r,it}, p_{r,it})'$; and income I_{it}. Given these variables, for each day t in bill period m, consumer i chooses her consumption vector – measured in kilowatt-hours (kWhs) – of off-peak and peak electricity $\boldsymbol{y}_{it} = (y_{op,it}, y_{p,it})'$ to maximize her value function,[8] shown in equation 1. The parameters η and λ (described in more detail below) are the parameters capturing reference-dependence and loss aversion, respectively. The first term of equation 1 is standard consumption utility over peak and off-peak electricity; the second term is utility over money (or the numeraire good); and the final bracketed term is the reference-dependent portion of utility.

$$
\begin{aligned}
(1)\quad U(\boldsymbol{y}_{it};\boldsymbol{x}_{it},p_{r,it}) &= u\left(\boldsymbol{y}_{it};\boldsymbol{x}_{it}\right) + \left(I_{it} - \boldsymbol{y}'_{it}\cdot\boldsymbol{p}_{it}\right) \\
&\quad - \begin{cases} \eta\lambda\left(\boldsymbol{y}'_{it}\cdot\boldsymbol{p}_{it} - \boldsymbol{y}'_{it}\cdot\boldsymbol{p}_{r,it}\right) & \text{if } \sum_{s\in m}\left[\boldsymbol{y}'_{is}\cdot\boldsymbol{p}_{is}\right] > \sum_{s\in m}\left[\boldsymbol{y}'_{is}\cdot\boldsymbol{p}_{r,is}\right] \\ \eta\left(\boldsymbol{y}'_{it}\cdot\boldsymbol{p}_{it} - \boldsymbol{y}'_{it}\cdot\boldsymbol{p}_{r,it}\right) & \text{if } \sum_{s\in m}\left[\boldsymbol{y}'_{is}\cdot\boldsymbol{p}_{is}\right] \le \sum_{s\in m}\left[\boldsymbol{y}'_{is}\cdot\boldsymbol{p}_{r,is}\right] \end{cases}
\end{aligned}
$$

In this model, the consumer has utility over bill period expenditure on electricity $\sum_{s\in m}\left[\boldsymbol{y}'_{is}\cdot\boldsymbol{p}_{is}\right]$ relative to a reference level of bill-period electricity expenditure, $\sum_{s\in m}\left[\boldsymbol{y}'_{is}\cdot\boldsymbol{p}_{r,is}\right]$, and incurs a loss if her current expenditure on electricity for the bill period is greater than her reference level of expenditure. The parameter η is the weight placed on the reference-dependent portion of utility relative to the direct consumption utility. It is assumed that $\eta \geq 0$, and $\eta = 0$ means the consumer has no reference-dependent utility. The loss-aversion parameter is λ; it is assumed that $\lambda \geq 1$, and if $\lambda = 1$ then the consumer is not loss averse – she cares equally about gains and losses relative

[8]"Value function" refers to consumption utility plus reference-dependent utility.

to her reference point – whereas $\lambda > 1$ means the consumer is loss averse, meaning losses relative to her reference point weigh more heavily in her utility than gains.

To make the model more realistic, I assume that the consumer is both imperfect at predicting her consumption on future days, and imperfect at recalling her consumption for past days.[9] Therefore, from the perspective of day t, I assume $\boldsymbol{y}_{is} = \hat{\boldsymbol{y}}_{is} - \boldsymbol{e}_{is}\ \forall s \neq t$, where $\boldsymbol{y}_{is}$ is consumer i's true observed consumption on day $s \neq t$; $\hat{\boldsymbol{y}}_{is}$ is her predicted or recalled consumption, and $\boldsymbol{e}_{is}$ is her prediction/recall error. This means that from the consumer's perspective, on any day t within a bill period, there is some probability she will experience a loss that bill period (δ_{it}), and her value function becomes that shown in equation 2.

$$\begin{aligned}
(2) \qquad U(\boldsymbol{y}_{it}; \boldsymbol{x}_{it}, p_{r,it}) &= u\left(\boldsymbol{y}_{it}; \boldsymbol{x}_{it}\right) + \left(I_{it} - \boldsymbol{y}'_{it} \cdot \boldsymbol{p}_{it}\right) \\
&\quad - \Delta_{it}\left(\boldsymbol{y}'_{it} \cdot \boldsymbol{p}_{it} - \boldsymbol{y}'_{it} \cdot \boldsymbol{p}_{r,it}\right)
\end{aligned}$$

where:

$$\begin{aligned}
\Delta_{it} &= \eta\lambda \cdot \delta_{it} + \eta \cdot (1 - \delta_{it}) \\
\delta_{it} &= prob_{it}\left(\sum_{s \in m}\left[\boldsymbol{y}'_{is} \cdot \boldsymbol{p}_{is}\right] > \sum_{s \in m}\left[\boldsymbol{y}'_{is} \cdot \boldsymbol{p}_{r,is}\right]\right) \\
&= prob_{it}\left(\sum_{s \in m} \hat{\boldsymbol{y}}'_{is} \cdot \left(\boldsymbol{p}_{is} - \boldsymbol{p}_{r,is}\right) > \sum_{s \in m} \boldsymbol{e}'_{is} \cdot \left(\boldsymbol{p}_{is} - \boldsymbol{p}_{r,is}\right)\right)
\end{aligned}$$

I am interested in the simple case where there has been an exogenous shock to the probability of incurring a loss (described in more detail in section III),

[9] This assumption does not drive the results of the model, only allows for imperfections in the degree to which consumers can predict their particular monthly expenditure and/or reference point.

so I assume δ_{it} is exogenous to the day t consumption decision.[10]

Because I make the common assumptions of quasi-linear utility, constant marginal utility of income,[11] and risk neutrality over both losses and gains in expenditure, the model has the convenient feature that the kink in the value function characterizing loss aversion is only present in the linear portion of the quasi-linear value function. This allows for a unified way, shown in equation 3, of representing the consumer's bill period value function for the two alternative models: the neoclassical model (no reference-dependence), and the reference-dependent model.

$$U(\boldsymbol{y}_{it}; \boldsymbol{x}_{it}, p_{r,it}) = u\left(\boldsymbol{y}_{it}; \boldsymbol{x}_{it}\right) + \left(I_{it} - \boldsymbol{y}'_{it} \cdot \bar{\boldsymbol{p}}_{it}\right) \tag{3}$$

where:

$$\bar{\boldsymbol{p}}_{it} = \begin{bmatrix} \bar{p}_{op,it} \\ \bar{p}_{p,it} \end{bmatrix} = \begin{bmatrix} p_{op,it} + \Delta_{it}\left(p_{op,it} - p_{r,it}\right) \\ p_{p,im} + \Delta_{it}\left(p_{p,it} - p_{r,it}\right) \end{bmatrix}$$

$$\Delta_{it} = \eta\lambda \cdot \delta_{it} + \eta \cdot (1 - \delta_{it})$$

In equation 3, $\bar{\boldsymbol{p}}_{it} = \left[\boldsymbol{p}_{it} + \left(\eta\lambda\delta_{it} + \eta\left(1 - \delta_{it}\right)\right)\left(\boldsymbol{p}_{it} - \boldsymbol{p}_{r,it}\right)\right]$ includes the reference-dependent features of the value function. If the consumer either has no reference-dependent utility (i.e., $\eta = 0$), or is a reference-dependent consumer on the standard (reference) pricing structure (i.e., $\boldsymbol{p}_{it} = \boldsymbol{p}_{r,it}$), then $\bar{\boldsymbol{p}}_{it}$ is simply her true prices, and the problem collapses to the standard problem with no reference dependence. However, if the consumer has reference-dependent

[10] Making δ_{it} endogenous to the daily choice of $\boldsymbol{y}_{it}$ would make the consumer's optimization problem more complicated, but would not qualitatively change the key predictions discussed here.

[11] Note that the assumption of constant marginal utility of income is not unreasonable, as the total expenditure on electricity is small relative to total income in general.

utility and is on the time-differentiated pricing structure, then $\bar{\boldsymbol{p}}_{it}$ reflects the way that reference-dependent utility differentially affects her electricity consumption, depending on her reference point, degree of loss aversion, and the probability she is in the loss domain.

Characterizing the problem in this way results in a specification of the value function that is continuous and everywhere twice differentiable in $\bar{\boldsymbol{p}}_{it}$ and $\boldsymbol{y}_{it}$, therefore any standard utility specification can be used for the consumption utility over electricity, $u(\boldsymbol{y}_{it}; \boldsymbol{x}_{it})$. Any duality properties in this model hold in terms of $\bar{\boldsymbol{p}}_{it}$, but not in terms of true prices, $\boldsymbol{p}_{it}$. In particular, Roy's Identity will hold with respect to $\bar{\boldsymbol{p}}_{it}$, but not with respect to $\boldsymbol{p}_{it}$.[12]

I.A Testable Prediction 1

The model predicts that the more likely it is that the consumer is in the loss domain with respect to bill period expenditure, the lower will be her peak consumption. To see this, note that each day t, consumer i's optimization problem is that shown in equation 4.

$$\max_{\boldsymbol{y}_{it}} U(\boldsymbol{y}_{it}; \boldsymbol{x}_{it}, p_{r,it}) = u(\boldsymbol{y}_{it}; \boldsymbol{x}_{it}) + (I_{it} - \boldsymbol{y}'_{it} \cdot \bar{\boldsymbol{p}}_{it}) \tag{4}$$

where:

$$\bar{\boldsymbol{p}}_{it} = \boldsymbol{p}_{it} + \Delta_{it} \left(\boldsymbol{p}_{it} - \boldsymbol{p}_{r,it} \right)$$

$$\Delta_{it} = \eta\lambda \cdot \delta_{it} + \eta \cdot (1 - \delta_{it})$$

The first order conditions for this optimization are shown in equation 5.[13]

[12] A proof of the fact that Roy's Identity holds with respect to $\bar{\boldsymbol{p}}_{it}$ is shown in online Appendix A.

[13] Recall that I assume here that $\frac{\partial \delta_{it}}{\partial \boldsymbol{y}_{it}} = 0$.

$$
\begin{cases} \frac{\partial u(\boldsymbol{y}_{it};\boldsymbol{x}_{it})}{\partial y_{op,it}} - \bar{p}_{op,it} & = 0 \\ \frac{\partial u(\boldsymbol{y}_{it};\boldsymbol{x}_{it})}{\partial y_{p,it}} - \bar{p}_{p,it} & = 0 \end{cases} \tag{5}
$$

The comparative static effect of an exogenous increase in the probability the consumer is in the loss domain for the bill period (δ_{it}) are shown in column 1 of table 1.[14]

Table 1: Comparative Static Effect of an Exogenous Increase in the Probability of a Loss

	Assuming $\frac{dy_{op}}{dp_p} = \frac{dy_p}{dp_{op}} > 0$ (1)	Assuming $\frac{dy_{op}}{dp_p} = \frac{dy_p}{dp_{op}} = 0$ (2)
$\frac{dy_{op}}{d\delta}$	$\frac{(\eta\lambda-\eta)}{(1+\Delta)}\left(\frac{dy_{op}}{dp_{op}}(p_{op}-p_r)+\frac{dy_{op}}{dp_p}(p_p-p_r)\right) > 0$	$\frac{(\eta\lambda-\eta)}{(1+\Delta)}\frac{dy_{op}}{dp_{op}}(p_{op}-p_r) \approx 0$
$\frac{dy_p}{d\delta}$	$\frac{(\eta\lambda-\eta)}{(1+\Delta)}\left(\frac{dy_p}{dp_{op}}(p_{op}-p_r)+\frac{dy_p}{dp_p}(p_p-p_r)\right) < 0$	$\frac{(\eta\lambda-\eta)}{(1+\Delta)}\frac{dy_p}{dp_p}(p_p-p_r) < 0$

Notes: This table shows the comparative static effect of increasing the probability the consumer is in the loss domain of her value function for the bill period, δ_{it}, on daily peak and off-peak electricity demand. Column 1 shows the full comparative static effect in the case of a loss-averse consumer, and column 2 shows this effect assuming that $\frac{dy_{op,it}}{dp_{p,it}} = \frac{dy_{p,it}}{dp_{op,it}} = 0$, and recognizing that both $\frac{dy_{op}}{dp_{op}}$ and $(p_{op} - p_r)$ are likely to be relatively small.

The comparative-static results indicate that $\frac{dy_{op,it}}{d\delta_{it}} > 0$ and $\frac{dy_{p,it}}{d\delta_{it}} < 0$ as long as $\eta > 0$, $\lambda > 1$, $p_{op,it} < p_{r,it}$, and $p_{p,it} > p_{r,it}$. This means the model predicts that, the higher the probability a consumer will experience a loss for the bill

[14] The full derivation of these comparative statics are shown in online Appendix B.

period, the more she will decrease peak consumption and increase off-peak consumption within that bill period in order to avoid a loss, *ceteris paribus*. However, I find very little empirical evidence in the data that off-peak usage is responsive to peak or off-peak price changes, meaning that the terms $\frac{dy_{op,it}}{dp_{p,it}}$ and $\frac{dy_{op,it}}{dp_{op,it}}$ are relatively small. Additionally, $(p_{op,it} - p_{r,it})$ is small relative to $(p_{p,it} - p_{r,it})$ in the data. Because of these reasons, the empirical predictions simplify down to those presented in column 2 of table 1.

In sum, the first testable prediction of the model is the following: if consumers are loss averse over bill period electricity expenditure, then the higher the probability that the consumer will incur a loss in a bill period, the lower their peak consumption will be during high-priced peak hours of that bill period in order to avoid a loss. If the consumer is not loss averse, then $\lambda = 1$, meaning there should be no correlation between exogenous changes in the probability the consumer is in the loss domain and the daily peak consumption behavior, because $\frac{\partial y_{p,it}}{\partial \delta_{it}} = 0$ in that case.

I.B Testable Prediction 2

The second testable prediction of the model is that if consumers are loss averse, there will be a disproportionate clustering of bill-period expenditure outcomes near where $\sum_{s \in m} [y_{op,is} \cdot p_{op,is} + y_{p,is} \cdot p_{p,is}] = \sum_{s \in m} [p_{r,is} \cdot (y_{op,is} + y_{p,is})]$ – the kink in the value function – particularly when households are in regions of their consumption that place them close to the kink or would otherwise skew them slightly into the loss domain. This section outlines the intuition for why this is the case.

Equation 1 demonstrates how the reference-dependent value function – before the introduction of uncertainty – is kinked. If there were no uncertainty, this kink in the value function would be sharp. Previous work, particularly in the area of labor supply, has posited clustering of outcomes at the kink to be a feature of kinked constraints. This would be equally true in the case of kinked objective functions. Moffitt (1990) provides a useful summary of this work, and discusses clustering at the kink in the budget constraint characterizing retirement age decisions found by Burtless and Moffitt (1984), and retirement consumption found by Burtless and Moffitt (1985). Another more recent example with empirical evidence for this kind of bunching at kink points, here in the context of tax schedule kink points, is provided by Saez (2010).

Because of the assumption that consumers are imperfect at controlling their electricity demand for the bill period, and therefore experience some uncertainty as to whether they will experience a loss or not, the kink in the case of this full model will be a "fuzzy kink," and could be thought of as more of a range of the value function that has extreme curvature. Outcomes would disproportionately cluster in this range of extreme curvature, similarly to how they would cluster at a sharp kink.

This clustering should be more pronounced if the relative prices are such that optimal outcomes are likely to be close to the fuzzy kink. It should particularly be more pronounced if the consumer's expenditure would otherwise be just skewed into the loss domain, because as we know from the first prediction of the model, the higher the probability the consumer is in the loss domain, the more she has an incentive is to cut back on expenditure by reducing peak consumption. This would cause her to pull back further towards the kink region.

The prediction of disproportionate clustering around the fuzzy kink can be tested, as the location of the kink is determined solely by observable prices. Therefore, the second testable prediction of this model is that, if consumers are loss averse over bill period expenditure on electricity, there is a disproportionate clustering of outcomes just around where $\sum_{s\in m}[y_{op,is}\cdot p_{op,is} + y_{p,is}\cdot p_{p,is}] = \sum_{s\in m}[p_{r,is}\cdot(y_{op,is} + y_{p,is})]$, particularly when households are in regions of their consumption that place their expenditure close to the kink and otherwise would skew their expenditure into the loss domain.

II Data

To test these predictions I use data from the California Statewide Pricing Pilot (SPP). This pilot was a collaboration between the California Energy Commission (CEC) and three of the state's largest electric utilities: Pacific Gas and Electric (PG&E); Southern California Edition (SCE); and San Diego Gas and Electric (SDG&E). The data consist of observations between roughly July 2003 and October 2004 of five groups: CPP High Ratio, CPP Low Ratio, TOU High Ratio, and TOU Low Ratio treatment groups, as well as a control group.[15] The control group was unaware that an experiment was being conducted, and were charged a standard time-invariant price for electricity. I use the term "reference price" to refer to the price control households faced, which is the same as the price treatment households within the same utility had been facing prior to the experiment, and would revert to if they dropped out of the experiment. I focus primarily on the two CPP treatments (described below), while using the TOU treatment groups and the control group as counterfactuals.

The two CPP treatment groups were charged a relatively high peak price

[15]Several different treatment groups were recruited for the pilot, but for this project I focus on a subset. In the terminology used in the original pilot, the treatment groups I used were the two CPP-F treatments and the two TOU treatments. I refer interested readers to previous analyses of this pilot for more detail on the other treatments (Herter 2007; Faruqui and George 2005; Charles River and Associates 2005).

for electricity from 2 pm to 7 pm on non-holiday weekdays, and a relatively low off-peak price. Additionally, the utility could call a limited number of critical peak days per season – announced the preceding day – wherein a precipitously high price was charged during peak hours. A total of twelve critical peak days were called during each of the two summer phases (May through October), and three were called during the one winter phase.[16] The choice to call a critical peak day depended on a variety of factors, including weather forecasts; system capacity and reliability; and the limit to the number of critical peak days that could be called. Utilities could only call critical peak days on non-holiday weekdays, but there was an attempt to call them on a variety of days of the week within that constraint. During the first summer of the experiment all critical peak days that were called were non-contiguous. In the second summer there were three sets of two or more proximate critical peak days. The two TOU treatment groups were also charged a relatively low off-peak price and a relatively high peak price (2 pm to 7 pm), though with no critical peak feature.[17] A table detailing all the experimental prices can be found in online Appendix C.

The two CPP treatments were called the CPP Low Ratio (CPPL) and CPP High Ratio (CPPH) treatments. The difference can be seen in figure 1, which plots the prices for non-CARE[18] PG&E customers in the two CPP treatments over the course of the pilot. The design of the prices was such that the CPPH treatment was expected to have relatively low expenditure in the

[16]This was true for PG&E and SDG&E, but SCE shifts from summer to winter pricing slightly earlier than the other two utilities, so three of the CPP days called that were in the summer of 2003 for the other two utilities were actually in the winter pricing phase for SCE.

[17]In the case of both the CPP and TOU tariffs, all consumption on weekends and holidays was charged at the off-peak price.

[18]CARE stands for California Alternate Rates for Energy, and is a program designed to provide price relief to low-income households.

summer and relatively high expenditure in the winter, while the opposite was true for the CPPL treatment. The critical peak prices are shown as points in the figures to demonstrate the frequency and timing of critical peak days.

Price data came from historic advice letters submitted by the utilities to the California Public Utilities Commission. California has an increasing block rate pricing structure for electricity. The time-based treatment prices consisted of a series of surcharges or credits overlaid onto this block rate structure, and were constant across the tiers. Because the theory of loss aversion used to motivate this analysis is primarily interested in prices relative to the reference price, the block rate structure is not directly relevant to my results. Additionally, previous research has shown that customers are not aware of, and do not respond to, the marginal price in their tiered structure, but rather respond to an averaged price (Ito 2014). I therefore conduct the analysis using the flat average price across the tiers. It is this average price that is plotted in figure 1 and reported in online Appendix C.

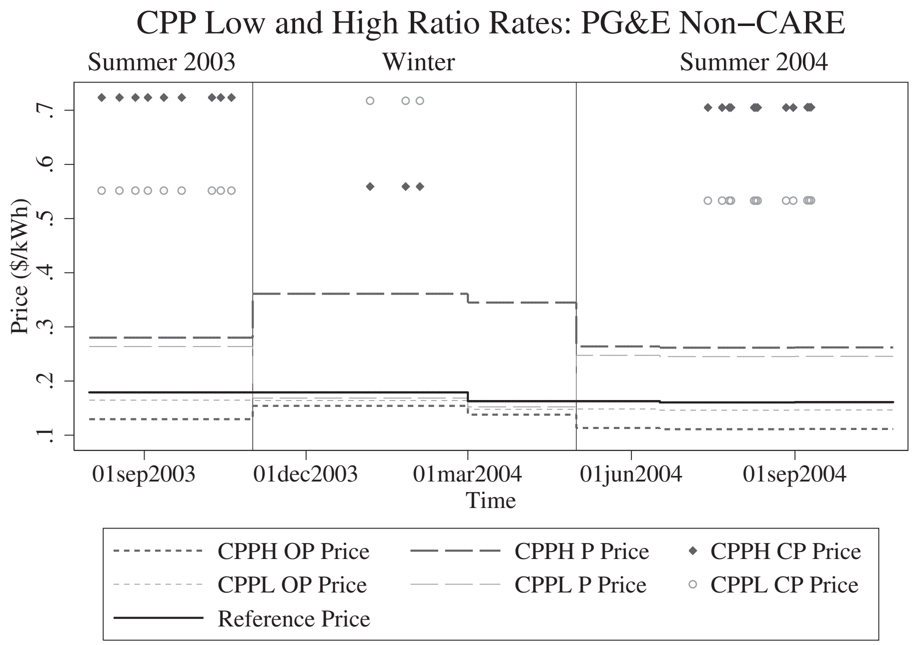

Figure 1: CPP Treatment Prices

Notes: This figure depicts the experimental prices faced by CPPH and CPPL non-CARE households in the PG&E service territory during the SPP experiment. "OP," "P," and "CP" refer to off-peak, peak, critical peak, respectively. "Reference Price" refers to the price charged to the control households. The frequency of critical peak pricing days is depicted in the points showing the critical peak prices. The prices shown are the average prices across the block rate tiers. The vertical lines represent the dates on which the pricing changed from summer to winter or vice versa.

The data include detailed electricity usage at 15-minute increments, which I identify as peak or off-peak usage, and aggregate up to the daily or bill period level. In addition, each household was matched to hourly temperature data from one of 56 weather stations. I construct a degree-hour measure of temperature within the peak and off-peak periods for each day. This measure is constructed similarly to the more commonly used degree-day measure, but separately for the peak and off-peak periods each day.[19]

[19]The degree-hour temperature measure is constructed in the following way: $dh_{op,t} =$

There are two main problems with the data from this experiment. First, there was some concern that households were unclear as to when precisely the experimental pricing started (Letzler 2010). I therefore drop observations from July 2003 (the initial month of the experiment). Second, the treatment groups were recruited to participate, while the control group was randomly selected from the population. This introduces an issue of selection into treatment. While I do use the control group as a comparison group, I run the same analyses using the TOU treatment groups as a counterfactual. This strengthens the comparability of the CPP and counterfactual groups, as they both selected into treatment.[20]

Table 2 shows summary statistics for the relevant variables. As one would expect, the treatment groups used less peak electricity on average than the control group during the experiment, and even at this aggregate level the difference is marginally significant. The difference in the off-peak electricity consumption between the four groups at this aggregate level is not statistically significant. This suggests broadly that the treatment did not induce a large amount of consumption shifting from peak to off-peak, the implication of which is that the own-price elasticity of off-peak, and cross-price elasticity between peak and off-peak, consumption are not large.

$|Mean\,(Temp_{h\in op,t}) - 65|$ and $dh_{p,t} = |Mean\,(Temp_{h\in p,t}) - 65|$, where $dh_{op,t}$ is the degree-hour temperature measure during off-peak hours on day t, $dh_{p,t}$ is the same measure but for peak hours on day t, $Mean\,(Temp_{h\in op,t})$ is average temperature during the off-peak hours of day t and $Mean\,(Temp_{h\in p,t})$ is the same for peak hours on day t.

[20] Online Appendix D outlines additional data cleaning determinations that were made, and presents robustness checks of the primary regressions reported in the paper to test the relevance of some of the data irregularities, none of which significantly change the results.

Table 2: Summary Statistics

	CPP High Ratio		CPP Low Ratio		TOU		Control	
	Mean	SD	Mean	SD	Mean	SD	Mean	SD
Average kWh per Day in 2002[a]	22.809	14.875	21.638	13.423	22.125	15.194	22.579	15.257
Off-Peak kWh per Day[b]	16.796	12.718	16.487	12.021	16.315	13.230	16.380	12.967
Peak kWh per Day[b]	5.362	5.388	5.396	5.494	5.468	5.823	6.158	6.588
Bill Total[b]	88.024	68.490	87.559	69.437	91.869	75.938	96.583	79.686
Off-Peak Degree-Hours[b]	8.691	6.361	8.611	6.287	8.908	6.195	8.917	6.292
Peak Degree-Hours[b]	12.442	9.256	12.267	9.247	12.074	9.427	12.418	9.678
PG&E Customer	0.485	0.500	0.489	0.500	0.624	0.485	0.547	0.498
SCE Customer	0.424	0.494	0.411	0.492	0.376	0.485	0.377	0.485
SDG&E Customer	0.091	0.287	0.101	0.301	0	0	0.076	0.265
Number of Observations	115109		118640		86590		152903	
Number of Households	321		345		240		418	

Notes: The two CPP treatments are the groups of interest in this study. The experimental TOU group and the control group are both used as counterfactuals. "Bill Total" is the average total bill-period expenditure, not including fixed charges or taxes, during the experiment. "SD" refers to standard deviation.
a Pretreatment
b During Treatment Period

Note that these four groups are comparable in terms of the pretreatment average daily usage from the summer of 2002, and in terms of temperature levels faced during treatment. However, the TOU group differs from the CPPH and CPPL groups based on other observables (as they were much more likely to come from the PG&E region, which means they are more likely to be from Northern California relative to the CPP groups). For this reason I use both the TOU and control groups as counterfactuals.

III Testing Model Predictions

In this section I present both the strategies I use to test the two predictions from the model, and the results from this analysis.

III.A Testable Prediction 1

Recall the first prediction from the model: The more likely it is the consumer is in the loss domain with respect to bill period expenditure for a given bill period, the more she will reduce her daily peak consumption within that bill period, *ceteris paribus*. To test this, I need an exogenous shock to the probability the consumer experiences a loss. In particular, I need an observable variable that is correlated with the probability of the household being in the loss domain for the bill period, but uncorrelated with electricity consumption on any given day.

Recall that I define a loss to be if the consumer paid more for her chosen electricity consumption in a given bill period than she would have *for the same level of consumption* on the standard pricing structure. Using a linear probability model, I regress an indicator variable – equal to one in the case a household incurred a loss during that bill period and zero otherwise – on the following variables: the share of the bill period that is considered "summer" in terms of the pricing structure (equal simply to zero or one for the majority of bill period observations); the number of critical peak days called in a given bill period; the average number of degree-hours in the peak and off-peak hours; and household fixed effects. I run this regression separately for the CPPH and CPPL households. The results for these regressions are presented in table 3.

Table 3: Linear Probability of Incurring a Bill Period Loss

Dependent Variable: Loss (0,1)	(1) CPP High Ratio	(2) CPP Low Ratio
Critical Peak Days (Number in Bill Period)	0.0350***	0.0773***
	(0.00395)	(0.00460)
Summer Pricing (Share of Bill Period)	-0.798***	0.400***
	(0.0256)	(0.0268)
Peak Temperature	0.00649***	0.00532***
	(0.00209)	(0.00165)
Off-Peak Temperature	0.00291	-0.00856***
	(0.00181)	(0.00150)
Constant	0.680***	0.0288
	(0.0258)	(0.0250)
Household fixed effects	Y	Y
Observations (bill periods)	4,067	4,194
Total Number of Households	321	345
R-squared (within)	0.549	0.538

Notes: The dependent variable is a bill-period-level indicator variable of whether or not the household incurred a loss that bill period. Standard errors clustered at household level are shown in the parentheses.
*** Significant at the 1 percent level.
** Significant at the 5 percent level.
* Significant at the 10 percent level.

The results in table 3 demonstrate that first, CPPH households were more likely to experience a loss in the winter pricing phase, while CPPL households were more likely to experience a loss in the summer pricing phase. This was intentional in the design of the experimental pricing. This is relevant, and will be discussed further in section III.B. Second, higher off-peak degree-hours had a significant negative effect on the probability of a loss only for the CPPL group. Third, the number of critical peak days experienced in a given bill period had a significant positive effect on the probability of a loss for both groups; for each additional critical peak day experienced in a bill period, the probability that the average household will experience a loss is increased by 3.5 and 7.73 percentage points for the CPPH and CPPL groups, respectively. Finally, the higher the number of degree-hours (a positive demand shock) in the high-priced peak hours, the more likely the households would incur a loss. The standard deviation for the peak degree-hour measure is between 9 and 10 for all groups. Therefore, according to these results, an increase in average peak degree-hours of one standard deviation would increase the probability of a loss by about 6.4 and 5.3 percentage points for the CPPH and CPPL groups, respectively.

Informed by the results from table 3, I define four variables for each day in the sample: (i) the number of critical peak days called so far in the bill period; (ii) the number of critical peak days called in the first week of the bill period; (iii) the average *peak* degree-hour temperature measure in the first week of the bill period; and (iv) the average *off-peak* degree-hour temperature measure in the first week of the bill period. I run regressions using each of the first three of these variables separately as exogenous positive shocks to the probability of a loss for the bill period. Additionally, I run one more regression using the last as a weak negative shock to the probability of a loss.

In all approaches I limit the analysis to the days in which peak hours are charged at peak prices (i.e., I eliminate weekends and holidays from the analysis). Additionally, to avoid issues of mean reversion, I limit the analysis to days at least one week after the shock to the probability of a loss occurred.[21] Finally, I compare the consumption of households in the treatment group to the consumption of a counterfactual group (either the control or TOU group)

[21] When using the number of critical peak days so far in a bill period, I limit the analysis to only days at least one week after the last critical peak day experienced, and for the three variables measuring critical peak days or degree-hours in the first week of a bill period, I limit the analysis to the third and fourth weeks of the bill period.

that did not experience critical peak pricing.

I use the estimating equation shown in equation 6. In this equation i denotes a household; t denotes a day; $y_{p,it}$, measured in kWhs, is daily peak electricity consumption; D_{it} is one of the four shocks to the probability of a loss defined above; $C_{it} \in \{0, 1\}$ is an indicator variable of whether or not the utility for household i called a critical peak day on day t; and $T_i \in \{0, 1\}$ is an indicator variable of whether or not household i is in one of the critical peak treatment groups. In the vector $\boldsymbol{x}_{p,it}$ I control for peak temperature measured in degree-hours on day t for household i; month-of-year effects; day-of-week effects; and whether day t for household i is in the summer or winter pricing phase. For the regressions using peak (off-peak) degree-hours in the first week of the bill period as a shock to the probability of a loss, I control for the off-peak (peak) degree-hours in the first week of the bill period as well. Finally, I include household fixed effects, γ_i. The parameters in the model are a, $\boldsymbol{d}$, and b_k, $k \in \{1, ..., 4\}$.

$$(6) \qquad y_{p,it} = a + \boldsymbol{d}' \cdot \boldsymbol{x}_{it} + b_1 C_{it} + b_2 D_{it} + b_3 T_i * C_{it} + b_4 T_i * D_{it} + \gamma_i + \varepsilon_{it}$$

The parameter of primary interest is the b_4 coefficient on $T_i * D_{it}$. Loss aversion would predict that this parameter be negative when D_{it} is a positive shock to the probability of a loss, meaning that the higher the probability of a loss due to previous expenditure shocks, the less peak electricity the household will consume subsequently in the bill period. In the case of a negative shock to the probability of a loss, b_4 would be expected to be positive.

Tables 4 presents the results from four versions of regressions based on equation 6.[22] Columns 1 and 2 show the results using the number of critical peak days so far in the bill period as a positive shock to the probability of a loss; columns 3 and 4 show the results using the number of critical peak days in the first week of the bill period as a positive shock to the probability of a loss; columns 5 and 6 show the results using average peak degree-hours in the first week of the bill period as a positive shock to the probability of a loss; and finally columns 7 and 8 show the results using average off-peak degree-hours in the first week of the bill period as a weak negative shock to the probability of a loss. Columns 1, 3, 5, and 7 present the results when the control group is used as the counterfactual, while columns 2, 4, 6, and 8 show results when the TOU group is used as the counterfactual.

The results in table 4 are relatively stable and comply with intuition: households consume more peak electricity when peak degree-hours are higher. Households consume more peak electricity on critical peak days, but treatment households respond to the higher critical peak prices and consume less (by 1.1 to 1.44 kWh relative to control households, which is about 20 percent of average daily peak electricity consumption) than control households on critical peak days.

The null hypotheses of no loss aversion are that the coefficients on $T_i * D_{it}$ equal zero. This null can be rejected in the first three sets of peak demand regressions relative to the control group. Looking at columns 1, 3, and 5 in table 4, consistently consumers reduce their peak consumption more relative to the control group if they have experienced a previous positive shock to the

[22] I run the same regressions with off-peak consumption as the dependent variable, the results of which can be found in online Appendix E. There are no significant loss-aversion results in these off-peak demand regressions, which is consistent with the model assuming little to no cross-price elasticity of off-peak consumption to peak prices (an assumption supported by the lack of significant impacts of critical peak prices on off-peak consumption).

probability of a loss. This is consistent with the model of loss aversion and not consistent with the standard neoclassical model. In particular, focusing first on column 1, an increase in the number of critical peak days experienced so far in a bill period by one day decreases the amount of peak electricity the household consumes subsequently in the bill period by 0.0914 kWhs per day (1.68 percent of average daily peak electricity consumption) on average for the CPP treatment groups relative to the control group. Column 3 shows that a one-day increase in the number of critical peak days experienced in the first week of the bill period resulted in a decrease in the peak electricity consumed on average in the last weeks of the bill period by 0.236 kWhs per day (4.37 percent of average daily peak consumption) for the CPP treatment groups relative to the control group. Finally, as shown in column 5, an increase of one degree-hour on average in the peak hours during the first week of the bill period resulted in a decrease in peak electricity consumption in the last weeks of the bill period of 0.0649 kWhs per day on average (1.20 percent) for the CPP treatment group relative to the control group.

When the average off-peak degree-hours in the first week of the bill period are used as a negative shock to the probability of a loss – results for which are presented in column 7 using the control group as the counterfactual – there is no significant affect on subsequent peak consumption during that bill period, though the sign of the coefficient is positive. This variable shows only a weak negative correlation with the probability of a loss, as it was only significant for the CPPL group in table 3. It is therefore not surprising that the effect is not significant here. However, it serves to support the validity of the results in columns 1, 3, and particularly 5, as it demonstrates that these results are not simply spurious correlations based on household behavior following what happened to have been hot days, for example.

Table 4: Peak kWh Usage Adjustment Following Shock to Loss Probability

Dependent Variable: Peak kWh	D = Number of critical peak days so far		D = Number of critical peak days in Week 1		D = Peak degree-hours in Week 1		D = Off-Peak degree-hours in Week 1	
	T=0: Control (1)	T=0: TOU (2)	T=0: Control (3)	T=0: TOU (4)	T=0: Control (5)	T=0: TOU (6)	T=0: Control (7)	T=0: TOU (8)
C	1.027***	0.442***	0.597***	0.300*	0.573***	0.365***	0.720***	0.430***
	(0.141)	(0.152)	(0.117)	(0.155)	(0.0989)	(0.131)	(0.118)	(0.156)
D	0.0512*	0.0704**	0.0584	0.0662	0.0920***	0.0696***	0.0351**	0.0210
	(0.0286)	(0.0354)	(0.0658)	(0.0797)	(0.0122)	(0.0147)	(0.0153)	(0.0191)
T * C	-1.438***	-0.690***	-1.296***	-0.789***	-1.112***	-0.771***	-1.358***	-0.862***
	(0.191)	(0.191)	(0.173)	(0.196)	(0.141)	(0.163)	(0.175)	(0.198)
T * D	-0.0914***	-0.0851**	-0.236***	-0.170**	-0.0649***	-0.0292	0.00640	0.0120
	(0.0345)	(0.0412)	(0.0729)	(0.0863)	(0.0167)	(0.0189)	(0.0145)	(0.0177)
Summer Pricing	-0.0896	-0.0894	0.118	0.0706	0.0780	0.0190	0.0769	0.0221
	(0.0743)	(0.0768)	(0.0833)	(0.0848)	(0.0840)	(0.0853)	(0.0836)	(0.0850)
Peak Degree-Hours	0.172***	0.155***	0.191***	0.170***	0.171***	0.153***	0.172***	0.153***
	(0.00748)	(0.00783)	(0.00803)	(0.00838)	(0.00684)	(0.00709)	(0.00687)	(0.00712)
Off-Peak Degree-Hours in Week 1					0.0369***	0.0297**		
					(0.0113)	(0.0122)		
Peak Degree-Hours in Week 1							0.0538***	0.0490***
							(0.00602)	(0.00601)
Constant	4.235***	3.996***	3.956***	3.758***	3.362***	3.256***	3.332***	3.235***
	(0.120)	(0.128)	(0.141)	(0.148)	(0.181)	(0.192)	(0.182)	(0.194)
Day-of-week effects	Y	Y	Y	Y	Y	Y	Y	Y
Month-of-year effects	Y	Y	Y	Y	Y	Y	Y	Y
Household fixed effects	Y	Y	Y	Y	Y	Y	Y	Y
Daily Observations	191,839	159,244	149,344	123,564	149,344	123,564	149,344	123,564
R-squared (within)	0.164	0.143	0.180	0.159	0.189	0.166	0.186	0.165
Average Peak kWh/day	5.45	5.12	5.4	5.1	5.4	5.1	5.4	5.1
Num. of Households (T=1)	666	666	655	655	655	655	655	655
Num. of Households (T=0)	418	240	416	237	416	237	416	237
Total Number of Households	1,084	906	1,071	892	1,071	892	1,071	892

Notes: The dependent variable is daily peak kWh usage. C is an indicator variable for the occurrence of a critical peak day; D is one of the four exogenous shocks to the probability of a loss; and T is an indicator variable equal to one for households in the pooled CPPL and CPPH treatment groups and equal to zero for households in the control and TOU counterfactual groups. Eight CPPH, three CPPL, two control, and three TOU households who appeared in the columns 1 and 2 regressions are dropped from the regressions presented in columns 3 through 8 because they have no observations in the third and fourth week of a bill period. A total of only 0.07 percent of observations are dropped due to the exclusion of these thirteen households. Standard errors clustered at the household level are shown in the parentheses.

*** Significant at the 1 percent level.
** Significant at the 5 percent level.
* Significant at the 10 percent level.

The same set of peak demand regressions using the TOU treatment group as the counterfactual can be seen in table 4 in columns 2, 4, 6, and 8. When the TOU group is used as the counterfactual the sign of the coefficient on $T_i * D_{it}$ is the same as when the control group is used as the counterfactual in all cases. However, this coefficient is only significantly different from zero in the two regressions using measures of previous critical peak days experienced shown in columns 2 and 4. The results indicate that an increase in the previous number of critical peak days experienced so far in the bill period by one day was associated with a reduction of 0.0851 kWh of peak consumption during subsequent peak hours relative to TOU households. This is 1.58 percent of average peak consumption. Additionally, a one-day increase in the number of critical peak days experienced in the first week of the bill period was associated with a reduction of 0.17 kWh (3.15 percent of average daily peak consumption) of peak electricity consumption relative to the TOU group. When shocks to the probability of a loss are in the form of higher peak degree-hours, the coefficient on $T_i * D_{it}$ is not significantly different from zero when the TOU group is the counterfactual. This makes sense. First, the TOU group is smaller than the control group, and so power is reduced. Second, the TOU group is also on an experimental pricing structure wherein higher prices are charged in peak hours and lower prices charged in off-peak hours. Therefore, shocks in the form of more critical peak days called would not increase the probability of a TOU household incurring a loss, but higher peak degree-hours would increase the probability of a TOU household themselves incurring a loss. Therefore, higher peak degree-hours in the first week of a TOU customer's bill period could also have induced her to cut back disproportionately on peak consumption later in the bill period in order to avoid her own loss. This is indicated by the fact that the coefficient on D_{it} is smaller for the TOU group in column 6 than for the control group in column 5.

This analysis demonstrates evidence that when there is a higher probability a household is in the loss domain due to previous positive exogenous shocks to bill period expenditure in the form of higher previous peak degree-hours, or more previous critical peak days, households cut back more on subsequent peak consumption.[23] This result is consistent with the model of loss aversion over electricity expenditure.

III.B Testable Prediction 2

Recall that section I.B outlined the model prediction that a disproportionate number of bill period expenditure outcomes should occur around the kink in the value function; particularly when prices are such that households tend to be otherwise located close to the fuzzy kink and skewed into the loss domain.

Recall that the kink is around where $\sum_{s\in m}\left[\boldsymbol{y}_{is}\cdot\boldsymbol{p}_{is}\right]=\sum_{s\in m}\left[\boldsymbol{y}_{is}\cdot\boldsymbol{p}_{r,is}\right]$. Note that this is equivalent to saying that expenditure net of reference expenditure is equal to zero ($\sum_{s\in m}\left[\boldsymbol{y}_{is}\cdot\boldsymbol{p}_{is}\right]-\sum_{s\in m}\left[\boldsymbol{y}_{is}\cdot\boldsymbol{p}_{r,is}\right]=0$). Both the CPPH and CPPL expenditure outcomes tend to be quite close to the kink throughout the experiment.[24] More compelling is that the CPPH treatment was designed such that households should be in the gain domain in the summer and be in the loss domain in the winter, whereas the CPPL treatment rates were designed so that they should be in the gain domain in the winter and the

[23] Robustness checks were run that omitted different subsets of the data. The results for these robustness checks are presented in online Appendix D, along with an explanation of why those checks were run. In all cases the results remain consistent with the primary regressions presented in table 4.

[24] The share of outcomes within \$13.2 (15 percent of the average bill) of the kink are 69 percent and 92 percent for the CPPH group during the summer and winter pricing phases, respectively. Conversely, for the CPPL treatment, 90 percent and 89 percent are in this range during the summer and winter pricing phases, respectively.

loss domain in the summer (Charles River and Associates 2005). The model predicts that we are more likely to see clustering in the season not only when households tend to be close to the kink, but also when they might otherwise have been skewed into the loss domain.

The treatment prices were constructed in such a way as to explicitly avoid large losses.[25] Therefore, clustering around the kink could occur by construction of the treatment tariffs themselves, and not because of household behavior. To ascertain whether the degree of clustering around the kink is disproportionately due to household behavior, rather than simply coincidental, I construct what I refer to as a counterfactual net expenditure by determining what control households would have spent if charged treatment prices minus what they actually spent having been charged control (reference) prices. Theoretically, the difference in the distribution of the treatment net expenditure (actual minus reference) and the counterfactual net expenditure should reflect behavioral change instigated by the treatment prices. I do the same process using the TOU households as the counterfactual as well.[26] These results are shown in figures 2 and 3.

In figures 2 and 3 the net-expenditure bins have widths of $8.8,[27] with the [-4.4,4.4] bin centered at $0 and considered to be "the kink." Figures 2 and 3 plot the mid-points of these bins on the horizontal axis, and the difference in share of bill-period net-expenditure observation in that bin between the treatment and counterfactual groups on the vertical axis. Results indicate that in winter

[25]The experimental rates had to meet three requirements: be revenue-neutral for the average customer over the year assuming unchanged load shape; not change the bills by more than 5 percent assuming unchanged load shape; and provide customers with the opportunity to save 10 percent on their bills if they reduced peak consumption by 30 percent (Charles River and Associates 2005).

[26]The TOU counterfactual net expenditure is TOU usage expenditure if charged CPP prices less TOU usage expenditure if charged control prices.

[27]10 percent of average bill totals for the CPP groups.

for the CPPH and summer for the CPPL groups there is clustering beyond that of the two counterfactuals right around the kink. These two phases are when the households were most likely to be located close to the kink, and otherwise skewed into the loss domain. The differences in the percent of net-expenditure outcomes for treatment households at the kink relative to the control counterfactual were 8.42 percent for CPPH in the winter and 11.36 percent for CPPL in the summer (2.14 percent for CPPH in the winter, and 4.82 percent for CPPL in the summer, relative to the TOU counterfactual). This pattern of clustering was not the case in the summer for the CPPH and winter for the CPPL groups, times in which these treatment groups were less likely to incur losses.

I test whether the degree of clustering observed for the CPPH group in the winter and CPPL group in the summer is statistically larger than the control and TOU counterfactuals by bootstrapping the distribution of the fraction of observations within the [-4.4,4.4] bin for each group.[28] A two-sample difference-in-means test, shown in table 5, demonstrates that indeed the probability of CPPL summer and CPPH winter outcomes being right around the kink are both statistically higher than the corresponding control and TOU counterfactual probabilities at the 99 percent confidence level.

[28] I drew with replacement a sample of households equal in number to the number of households in the data. I do this within each treatment and pricing phase separately. I repeated this process 2,000 times, recording the share of total observations with net-expenditure outcomes in the range [-4.4,4.4] for each case for all 2,000 repetitions. The resulting set of 2,000 data-points in each case made up the distribution of the probability of the share of net-expenditure observations at the kink in each case, clustered at the household level.

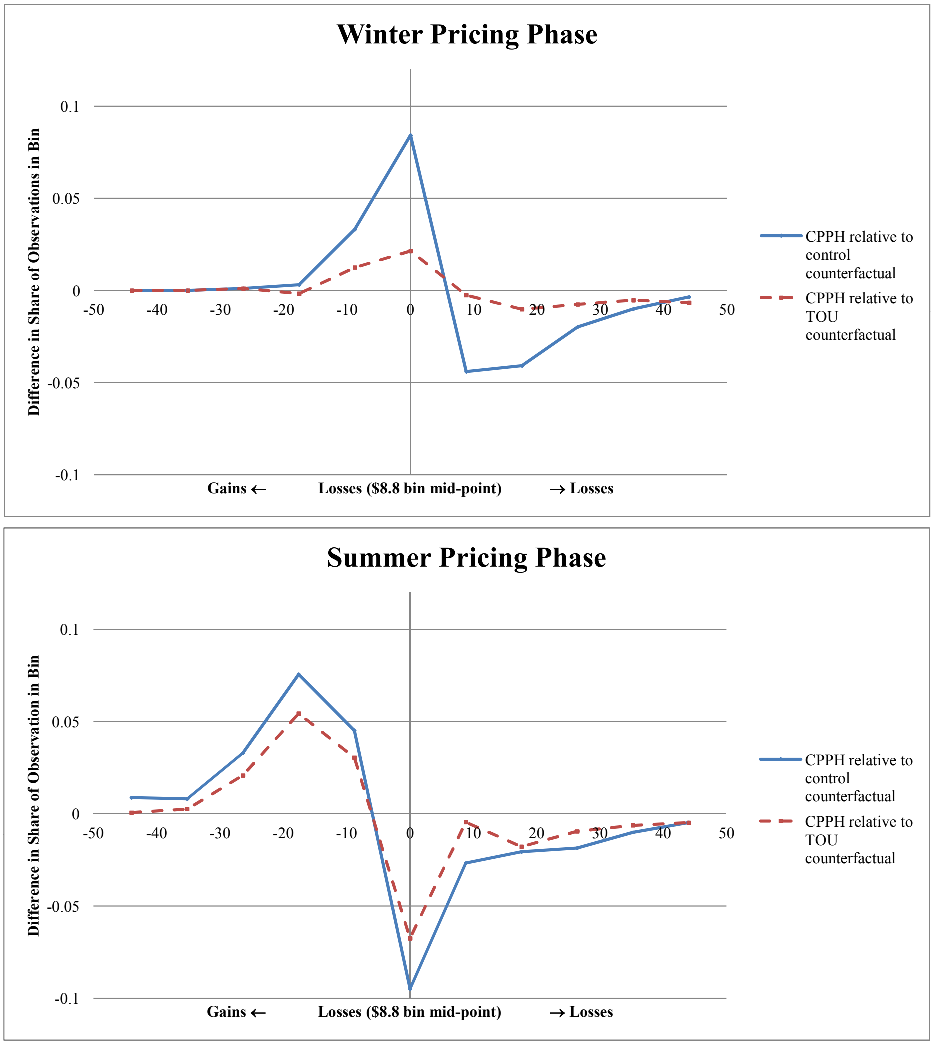

Figure 2: CPPH Bill Period Expenditure Clustering

Notes: This figure shows CPPH results for clustering around the kink. Both the control (solid line) and TOU (dashed line) households are used as counterfactuals. The center-points of bins with widths of \$8.8 representing net-expenditure ($\sum_{s \in m} [\boldsymbol{y}_{is} \cdot \boldsymbol{p}_{is}] - \sum_{s \in m} [\boldsymbol{y}_{is} \cdot \boldsymbol{p}_{r,is}]$) are plotted on the horizontal axis. The share of treatment bill-period net-expenditure observation in that bin minus the share of counterfactual bill-period net-expenditure observations in that bin are plotted on the vertical axis. The bin centered at zero and including the range [-4.4,4.4] is considered to be the "fuzzy kink" in the value function. Positive net-expenditure values are in the loss domain, while negative net-expenditure values are in the gain domain. The top panel shows bill periods in which more than 50 percent of the days were charged at winter prices, while the bottom panel shows bill periods that have more than 50 percent of days were charged at summer prices.

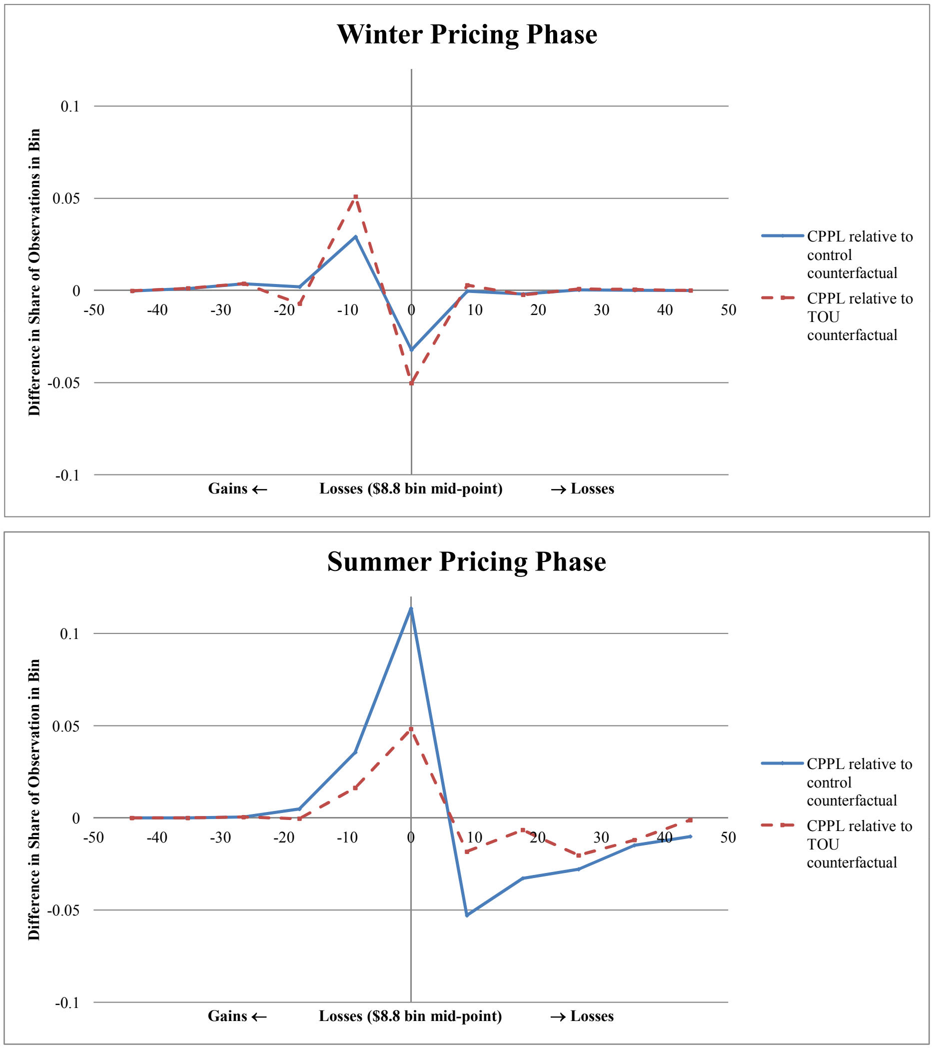

Figure 3: CPPL Bill Period Expenditure Clustering

Notes: This figure shows CPPL results for clustering around the kink. Both the control (solid line) and TOU (dashed line) households are used as counterfactuals. The center-points of bins with widths of $8.8 representing net-expenditure ($\sum_{s \in m} [\boldsymbol{y}_{is} \cdot \boldsymbol{p}_{is}] - \sum_{s \in m} [\boldsymbol{y}_{is} \cdot \boldsymbol{p}_{r,is}]$) are plotted on the horizontal axis. The share of treatment bill-period net-expenditure observation in that bin minus the share of counterfactual bill-period net-expenditure observations in that bin are plotted on the vertical axis. The bin centered at zero and including the range [-4.4,4.4] is considered to be the "fuzzy kink" in the value function. Positive net-expenditure values are in the loss domain, while negative net-expenditure values are in the gain domain. The top panel shows bill periods in which more than 50 percent of the days were charged at winter prices, while the bottom panel shows bill periods that have more than 50 percent of days were charged at summer prices.

Table 5: Bootstrapped Distributions of Share of Outcome at the Kink

	Treatment		Control		TOU		Treatment vs. Control	Treatment vs. TOU
	Mean	SD	Mean	SD	Mean	SD	t-stat	t-stat
CPPH Winter	0.59	0.02	0.51	0.02	0.57	0.03	51.52***	9.51***
CPPL Summer	0.69	0.02	0.58	0.02	0.65	0.02	92.25***	30.19***

Notes: Bootstrapped mean and standard deviations of the probability that a net-expenditure outcome will be within $4.4 of zero for the CPPH treatment in the winter pricing phase and the CPPL treatment in the summer pricing phase. The standard deviations are clustered at the household level by construction in the bootstrapping process, and significance is based on a t-test of the difference in means.
*** Significant at the 1 percent level.
** Significant at the 5 percent level.
* Significant at the 10 percent level.

One might suggest that demonstrating clustering in this way would only indicate that the curvature of the utility function was extreme, but would not necessarily imply loss aversion. However, for clustering to result near where $\sum_{s\in m}[\boldsymbol{y}_{is}\cdot\boldsymbol{p}_{is}]-\sum_{s\in m}\left[\boldsymbol{y}_{is}\cdot\boldsymbol{p}_{r,is}\right]=0$ from simply an extremely curved utility function, one would have to assume that the location of the extremely curved ridge in the utility function be the same for everyone, and happen to correspond to where expenditure (given differing prices and demand shifters) was equal to what it would have been on reference prices. There is no immediate reason, beyond the loss aversion model proposed in this paper, why this should be the case. Therefore, there is evidence that treatment households did exhibit statistically significant disproportionate clustering around the kink with respect to bill period expenditure during phases that were more likely to place them close to the kink and otherwise skew them into the loss domain.

IV Alternate Hypotheses

In the preceding sections I have demonstrated that consumers on the CPP treatments in the SPP experiment exhibited behavior consistent with a model of loss aversion over electricity expenditure. In this section I test two alternative hypotheses regarding the disproportionate reduction in peak electricity consumption following positive expenditure shocks in a given bill period.

IV.A Learning Strategies to Reduce Peak Consumption

In section III.A I showed that when households on the CPP experimental tariffs have experienced more critical peak days in a given bill period, they disproportionately cut back on subsequent peak electricity consumption. This behavior is consistent with the model of loss aversion over electricity expenditure presented in this paper. However, if households learned better strategies for reducing peak electricity consumption after experiencing more critical peak days, this could also explain the empirical results presented above. To explore this question, I present the results from two alternative specifications in table 6: In the first, I control simply for the total number of critical peak days experienced so far in the pilot overall, and in the second I allow this variable to enter quadratically.[29]

[29]The results presented here are for the case with the control group as the counterfactual. I provide the results using the TOU group as the counterfactual in online Appendix F.

Table 6: Learning vs. Loss Aversion: CPP vs. Control

T=1: CPP T=0: Control Dependent Variable: Peak kWh	E = overall number of critical days so far					
	D = Number of critical peak days so far		D = Number of critical peak days in Week 1		D = Peak degree-hours in Week 1	
	(1)	(2)	(3)	(4)	(5)	(6)
C	1.014***	0.992***	0.601***	0.544***	0.578***	0.555***
	(0.142)	(0.140)	(0.116)	(0.109)	(0.0976)	(0.0959)
D	0.0614**	0.0346	0.0701	-0.00904	0.0926***	0.0878***
	(0.0265)	(0.0259)	(0.0628)	(0.0648)	(0.0122)	(0.0126)
T * C	-1.438***	-1.400***	-1.293***	-1.201***	-1.111***	-1.074***
	(0.192)	(0.187)	(0.170)	(0.159)	(0.139)	(0.135)
T * D	-0.0878***	-0.0378	-0.232***	-0.0982	-0.0652***	-0.0571***
	(0.0327)	(0.0307)	(0.0696)	(0.0706)	(0.0166)	(0.0176)
E	-0.00809	-0.0949***	-0.00682	-0.135***	-0.00982	-0.0794**
	(0.00881)	(0.0259)	(0.0102)	(0.0381)	(0.0104)	(0.0367)
E^2		0.00340***		0.00478***		0.00258**
		(0.000991)		(0.00134)		(0.00126)
T * E	-0.00189	0.114***	-0.00153	0.183***	0.00179	0.0986**
	(0.0105)	(0.0320)	(0.0125)	(0.0447)	(0.0124)	(0.0435)
T * E^2		-0.00458***		-0.00702***		-0.00366**
		(0.00119)		(0.00163)		(0.00154)
Summer Pricing	-0.0917	-0.0789	0.116	0.140*	0.0779	0.0901
	(0.0742)	(0.0743)	(0.0831)	(0.0828)	(0.0839)	(0.0837)
Peak Degree-Hours	0.172***	0.172***	0.191***	0.191***	0.171***	0.172***
	(0.00748)	(0.00749)	(0.00802)	(0.00802)	(0.00684)	(0.00685)
Off-Peak Degree-Hours					0.0356***	0.0354***
in Week 1					(0.0114)	(0.0114)
Constant	4.345***	4.423***	4.051***	4.146***	3.472***	3.530***
	(0.126)	(0.154)	(0.145)	(0.215)	(0.185)	(0.247)
Day-of-week effects	Y	Y	Y	Y	Y	Y
Month-of-year effects	Y	Y	Y	Y	Y	Y
Household fixed effects	Y	Y	Y	Y	Y	Y
Daily Observations	191,839	191,839	149,344	149,344	149,344	149,344
R-squared (within)	0.164	0.164	0.180	0.181	0.189	0.190
Number of Households	1,084	1,084	1,071	1,071	1,071	1,071

Notes: The dependent variable is daily peak kWh usage. C is an indicator variable for the occurrence of a critical peak day; D is one of the four exogenous shocks to the probability of a loss; and T is an indicator variable equal to one for households in the pooled CPPL and CPPH treatment groups and equal to zero for households in the control counterfactual group. Columns 1, 3, and 5 control for the total number of critical peak days experienced in the pilot overall (E and $T * E$), and columns 2, 4, and 6 allow these variables to enter quadratically. Standard errors clustered at household level are shown in the parentheses.
*** Significant at the 1 percent level.
** Significant at the 5 percent level.
* Significant at the 10 percent level.

The results from this analysis are quite interesting. First, when the overall past experience of critical peak days for the treatment group ($T * E$) enters linearly – shown in columns 1, 3, and 5 of table 6 – the coefficients on this term are never significant, and additionally there is almost no difference in the coefficient on $T*D$ compared to the original specification shown in table 4. When the experience of overall past critical peak days enters quadratically – shown in columns 2, 4, and 6 – the coefficient on $T * E$ is significantly *greater than zero* in all specifications. This means that, rather than more experience of critical peak days overall resulting in increasingly *less* peak consumption, which would be consistent with learning, these results indicate that more critical peak days experienced overall is correlated with a subsequent relative *increase* in peak consumption. This is not consistent with learning better strategies to reduce peak consumption. What then could explain these results?

Köszegi and Rabin (2006) suggest that consumers do not have a stagnant reference point, but rather update their reference point based on expectations formed by recent experience. In the SPP pilot setting, if households became more accustomed to expenditure fluctuations as a result of critical peak days, then they may update their reference expenditure to reflect this. This updating of the reference point would explain a gradual relaxing of the degree to which households mitigate peak electricity consumption when they are exposed to more critical peak days.

The coefficient on $T*D$ is no longer significant in columns 2 and 4, possibly due to correlation with E. However, in column 6 higher peak degree-hours in the first week of the bill period still significantly correlates with lower subsequent peak consumption in the bill period relative to the control group. This is consistent with households updating their reference point with respect to experiencing critical peak days, but still responding when there is a higher probability of a loss, even given an updated reference point, as a result of exogenous weather shocks.

This analysis suggests that learning better strategies to reduce peak consumption cannot explain the original results that households cut back on peak consumption after experiencing more critical peak days in a bill period. Rather, the results indicate that households may actually *reduce* the degree to which they are curbing peak consumption the more they get used to being charged higher peak prices. This is still consistent with households having loss-averse preferences, but suggests that they may be updating their reference point to reflect more familiarity with the CPP pricing tariff.

IV.B Constrained Budget

In this section I address the possible alternative explanation that households cut back on peak electricity consumption after previous positive shocks to expenditure in the bill period not because they are loss averse, but rather because they would experience high penalties if they spend more than they are expecting to on electricity.

I split the households into CARE and non-CARE households. The latter group I further subdivide by income level. The majority of CARE households – which faced lower prices than non-CARE households – have household incomes less than $25,000 annually in my sample. Income levels are not observed for all households in the data.[30] Of non-CARE households 83 and 85 percent of CPP and TOU households, respectively, completed the income portion of the survey, as did 65 percent of control households. Just over 60 percent non-CARE households made between $25,000 and $100,000 per year.

[30]The income variable was obtained through a survey that was not returned by all participants.

Table 7: Budget Constrained vs. Loss Averse: CPP vs. Control

T=1: CPP T=0: Control Dependent Variable: Peak kWh	D = Number of critical peak days so far (1)		D = Number of critical peak days in Week 1 (2)		D = Peak degree-hours in Week 1 (3)	
C * (CARE)	0.392	(0.408)	0.166	(0.416)	0.305	(0.338)
C * (Income<25,000)	0.566	(0.735)	0.171	(0.527)	0.200	(0.362)
C * (25,000<Income<50,000)	0.627*	(0.357)	0.330	(0.344)	0.449	(0.292)
C * (50,000<Income<75,000)	1.477***	(0.381)	1.242***	(0.368)	1.095***	(0.304)
C * (75,000<Income<100,000)	1.471***	(0.516)	1.372**	(0.600)	1.290***	(0.469)
C * (100,000<Income<150,000)	1.796***	(0.667)	1.085**	(0.484)	1.061**	(0.421)
C * (150,000<Income)	1.683***	(0.634)	0.667	(0.452)	0.622	(0.410)
D * (CARE)	0.0461	(0.0797)	0.184	(0.180)	0.0397	(0.0324)
D * (Income<25,000)	0.0379	(0.0983)	-0.0907	(0.176)	0.0682	(0.0618)
D * (25,000<Income<50,000)	-0.0289	(0.0543)	-0.0692	(0.125)	0.0338*	(0.0188)
D * (50,000<Income<75,000)	0.122*	(0.0667)	0.283*	(0.170)	0.127***	(0.0305)
D * (75,000<Income<100,000)	0.290***	(0.0991)	0.336**	(0.168)	0.194***	(0.0490)
D * (100,000<Income<150,000)	0.0486	(0.0893)	-0.0932	(0.152)	0.143***	(0.0406)
D * (150,000<Income)	0.0832	(0.121)	-0.0987	(0.235)	0.122***	(0.0294)
T * C * (CARE)	-0.551	(0.455)	-0.671	(0.457)	-0.694*	(0.370)
T * C * (Income<25,000)	-1.753**	(0.777)	-1.303**	(0.604)	-0.974**	(0.447)
T * C * (25,000<Income<50,000)	-1.418***	(0.396)	-1.233***	(0.394)	-1.097***	(0.337)
T * C * (50,000<Income<75,000)	-1.886***	(0.497)	-1.995***	(0.496)	-1.727***	(0.398)
T * C * (75,000<Income<100,000)	-1.895***	(0.568)	-1.984***	(0.659)	-1.691***	(0.524)
T * C * (100,000<Income<150,000)	-1.415*	(0.774)	-1.339**	(0.587)	-1.518***	(0.514)
T * C * (150,000<Income)	-2.349***	(0.701)	-1.679***	(0.565)	-1.497***	(0.514)
T * D * (CARE)	-0.0985	(0.0891)	-0.348*	(0.199)	-0.0119	(0.0359)
T * D * (Income<25,000)	-0.0931	(0.114)	-0.111	(0.214)	-0.101	(0.0659)
T * D * (25,000<Income<50,000)	-0.0309	(0.0678)	-0.196	(0.144)	-0.0393	(0.0248)
T * D * (50,000<Income<75,000)	-0.120	(0.0876)	-0.443**	(0.206)	-0.0902**	(0.0439)
T * D * (75,000<Income<100,000)	-0.324***	(0.112)	-0.498***	(0.192)	-0.177***	(0.0543)
T * D * (100,000<Income<150,000)	0.0553	(0.118)	0.255	(0.225)	0.00108	(0.0525)
T * D * (150,000<Income)	-0.292**	(0.139)	-0.359	(0.265)	-0.0852**	(0.0423)
Summer Pricing	-0.117	(0.0766)	0.0594	(0.0898)	0.0204	(0.0916)
Peak Degree-Hours	0.172***	(0.00783)	0.189***	(0.00833)	0.170***	(0.00713)
Off-Peak Degree-Hours in Week 1					0.0399***	(0.0112)
Constant	4.222***	(0.127)	3.994***	(0.152)	3.378***	(0.188)
Day-of-week effects	Y		Y		Y	
Month-of-year effects	Y		Y		Y	
Household fixed effects	Y		Y		Y	
Daily Observations	163,035		127,007		127,007	
R-squared (within)	0.167		0.181		0.196	
Total Number of Households	872		863		863	

Notes: The dependent variable is daily peak kWh usage. C is an indicator variable for the occurrence of a critical peak day; D is one of the four exogenous shocks to the probability of a loss; and T is an indicator variable equal to one for households in the pooled CPPL and CPPH treatment groups and equal to zero for households in the control counterfactual group. The analysis is differentiates results between CARE households, as well as six income brackets of non-CARE households that provided answers to the income question in the survey. Standard errors clustered at household level are shown in the parentheses.
*** Significant at the 1 percent level.
** Significant at the 5 percent level.
* Significant at the 10 percent level.

Table 7 presents the results differentiated across CARE status, as well as income levels for the non-CARE households.[31] While the coefficient on $T * D$ is still negative for CARE households in all specifications, it is only marginally significant in the specification in column 2. It does not appear that CARE households were a strong driver of the loss aversion effect. Rather, depending on the specification, households in the \$50,000 to \$75,000; \$75,000 to \$100,000; and over \$150,000 annual income brackets appear to drive the result; households in these three mid- to high-income brackets significantly cut back on peak consumption following previous positive shocks to electricity expenditure in that bill period.

These results do not support the hypothesis that this behavior is being driven by households with a binding bill period budget constraint that would cause them to face high penalties if they overspent on electricity, as electricity expenditure for the households exhibiting this behavior constituted at most around 2 percent of monthly household income. These results therefore support the hypothesis that this pattern is due to behavioral factors, such as loss aversion, and not because households need to avoid overspending on electricity because of other constraints.

V Conclusion

In this study I specify a model of loss aversion over electricity expenditure. I use this model as a basis to derive testable predictions for consumption behavior of households on two critical peak pricing experimental tariffs. I use exogenous variation in the number of critical peak days called and the peak-

[31]The control group is used as the counterfactual. Results with the TOU group as the counterfactual can be seen in online Appendix F.

period degree-hours in the bill period, which are positively correlated with the probability of ending up in the loss domain of reference-dependent utility; when households are more likely to be in the loss domain, they respond by more strongly cutting back on peak consumption during subsequent peak hours. The magnitude of the additional reduction in peak consumption relative to the counterfactual group was between 1.20 to 4.37 percent of daily peak electricity consumption, depending on the positive shock to the probability of a loss and the counterfactual group used. Additionally, I show evidence of disproportionate clustering of bill-period expenditure around the kink in the reference-dependent value function during pricing phases that placed households close to the kink, and particularly that otherwise would have skewed them into the loss domain.

Finally, I explore two alternative explanations for this behavior other than loss aversion. First, I show that the reduction in peak consumption following positive exogenous shocks to electricity expenditure does not appear to be a result of households learning new strategies for reducing peak consumption when they experienced more critical peak days overall. Indeed, treatment households actually appeared to increase their peak consumption, relatively speaking, at a decreasing rate, the more critical peak days they experienced overall during the pilot. This behavior is not only counter to a learning explanation, but is consistent with households updating their reference point over time. Second, I show that households exhibiting the strongest evidence of reducing peak consumption following positive expenditure shocks tended to be mid- to high-income households, for whom electricity expenditure only constituted at most around 2 percent of monthly income, indicating that this behavior does not appear to be explained by households being severely budget constrained such that they would incur a high penalty if they overspent on electricity.

In essence these results demonstrate that the occurrence of critical peak days did not only result in a reduction of peak consumption on that day, but also spilled over to further reduction of peak consumption on regular peak days for several weeks thereafter. This was similarly true when degree-hours were high during high-priced periods. This form of demand adjustment resulted in households experiencing bill-period expenditures equal to what they would have paid on the standard time-invariant pricing tariff at a disproportionate rate. This higher number of bill periods with equal expenditure displaced bill periods in which they otherwise would have paid more than if they were on standard pricing.

The results from this analysis are relevant for the design and implementation of time-based pricing structures for several reasons. First, the documented greater effectiveness of CPP as compared to TOU or PTR structures found in a number of pilots is explained somewhat by the presence of loss aversion, as the intermittence of significantly higher prices called during critical peak days on CPP tariffs creates variation in expenditure in the loss domain of consumers' reference-dependent value function. This underscores the value of a pricing structure such as CPP from the perspective of long-term efficiency and sustainability of the electricity grid, as it provides a theoretical underpinning to support the empirical results found more generally that CPP rates are more effective at reducing consumption during the highest peak demand periods. Second, while CPP tariffs have been shown to be widely effective at reducing peak consumption, there is concern that consumers dislike them, and so would resist CPP tariffs. This dislike may stem from, in part, the comparison between the status-quo reference expenditure and the expenditure on a CPP pricing structure. A loss-averse consumer would be more likely to dislike a CPP pricing structure because of the nature of the variation in expenditure on these structures. However, policies that would allow consumers to ease into the experience of a CPP structure, or ones that would update their reference-point in advance by providing "shadow" bills showing how much consumers on the status-quo tariff would be paying on the CPP tariff, would be expected to reduce the resistance to a CPP tariff.

References

Allcott, Hunt. 2011. "Social norms and energy conservation." *Journal of Public Economics*, 95(9): 1082–1095.

Allcott, Hunt, and Todd Rogers. 2014. "The Short-Run and Long-Run Effects of Behavioral Interventions: Experimental Evidence from Energy Conservation." *American Economic Review*, 104(10): 3003–3037.

Arkes, Hal R, David Hirshleifer, Danling Jiang, and Sonya Lim. 2008. "Reference point adaptation: Tests in the domain of security trading." *Organizational Behavior and Human Decision Processes*, 105(1): 67–81.

Ayres, Ian, Sophie Raseman, and Alice Shih. 2009. "Evidence from two large field experiments that peer comparison feedback can reduce residential energy usage." National Bureau of Economic Research Working Paper 15386.

Barberis, Nicholas, and Wei Xiong. 2009. "What drives the disposition effect? An analysis of a long-standing preference-based explanation." *The Journal of Finance*, 64(2): 751–784.

Barberis, Nicholas, Ming Huang, and Tano Santos. 2001. "Prospect theory and asset prices." *The Quarterly Journal of Economics*, 116(1): 1–53.

Baucells, Manel, Martin Weber, and Frank Welfens. 2011. "Reference-point formation and updating." *Management Science*, 57(3): 506–519.

Benartzi, Shlomo, and Richard H Thaler. 1995. "Myopic loss aversion and the equity premium puzzle." *The Quarterly Journal of Economics*, 110(1): 73–92.

Borenstein, Severin, Michael Jaske, and Arthur Rosenfeld. 2002. "Dynamic pricing, advanced metering, and demand response in electricity markets." Center for the Study of Energy Markets Working Paper 105.

Burtless, Gary, and Robert Moffitt. 1984. "The effect of social security benefits on the labor supply of the aged." In *Retirement and Economic Behavior.* , ed. Henry Aaron and Gary Burtless, 135–171. Washington, DC: Brookings Institution.

Burtless, Gary, and Robert Moffitt. 1985. "The joint choice of retirement age and postretirement hours of work." *Journal of Labor Economics*, 3(2): 209–236.

Camerer, Colin F. 2000. "Prospect theory in the wild: Evidence from the field." In *Choices, Values, and Frames.* , ed. Daniel Kahneman and Amos Tversky, 288–300. New York: Cambridge University Press.

Camerer, Colin, Linda Babcock, George Loewenstein, and Richard Thaler. 1997. "Labor supply of New York City cab drivers: One day at a time." *The Quarterly Journal of Economics*, 112(2): 407–441.

Costa, Dora, and Matthew Kahn. 2013. "Energy conservation "nudges" and environmentalist ideology: Evidence from a randomized residential electricity field experiment." *Journal of the European Economic Association*, 11(3): 680–702.

CRA. 2005. "Impact Evaluation of the California Statewide Pricing Pilot." Charles River and Associates. Unpublished Technical Report.

De Giorgi, Enrico G, and Thierry Post. 2011. "Loss aversion with a state-dependent reference point." *Management Science*, 57(6): 1094–1110.

DellaVigna, Stefano. 2009. "Psychology and economics: Evidence from the field." *Journal of Economic Literature*, 47(2): 315–372.

Farber, Henry S. 2005. "Is tomorrow another day? The labor supply of New York City cabdrivers." *Journal of Political Economy*, 113(1): 46–82.

Farber, Henry S. 2008. "Reference-dependent preferences and labor supply: The case of New York City taxi drivers." *The American Economic Review*, 98(3): 1069–1082.

Faruqui, Ahmad. 2010. "Dynamic Pricing and Customer Behavior." *Brattle Group Presentation.*

Faruqui, Ahmad, and Stephen George. 2005. "Quantifying customer response to dynamic pricing." *The Electricity Journal*, 18(4): 53–63.

Fehr, Ernst, and Lorenz Goette. 2007. "Do workers work more if wages are high? Evidence from a randomized field experiment." *The American Economic Review*, 97(1): 298–317.

Genesove, David, and Christopher Mayer. 2001. "Loss aversion and seller behavior: Evidence from the housing market." *The Quarterly Journal of Economics*, 116(4): 1233–1260.

Hardie, Bruce GS, Eric J Johnson, and Peter S Fader. 1993. "Modeling loss aversion and reference dependence effects on brand choice." *Marketing Science*, 12(4): 378–394.

Hartman, Raymond S, Michael J Doane, and Chi-Keung Woo. 1991. "Consumer rationality and the status quo." *The Quarterly Journal of Economics*, 106(1): 141–162.

Herter, Karen. 2007. "Residential implementation of critical-peak pricing of electricity." *Energy Policy*, 35(4): 2121–2130.

Ito, Koichiro. 2014. "Do Consumers Respond to Marginal or Average Price? Evidence from Nonlinear Electricity Pricing." *American Economic Review*, 104(2): 537–63.

Jessoe, Katrina, and David Rapson. 2014. "Knowledge is (less) power: Experimental evidence from residential energy use." *American Economic Review*, 104(4): 1417–1438.

Joskow, Paul, and Catherine Wolfram. 2012. "Dynamic Pricing of Electricity." *American Economic Review: Papers & Proceedings*, 102(3): 381–385.

Kahneman, Daniel, and Amos Tversky. 1979. "Prospect theory: An analysis of decision under risk." *Econometrica: Journal of the Econometric Society*, 47(2): 263–291.

Köszegi, Botond, and Matthew Rabin. 2006. "A model of reference-dependent preferences." *The Quarterly Journal of Economics*, 121(4): 1133–1165.

Letzler, Robert. 2010. "Using incentive preserving rebates to increase acceptance of critical peak electricity pricing." Center for the Study of Energy Markets Working Paper 162.

Li, Yan, and Liyan Yang. 2013. "Prospect theory, the disposition effect, and asset prices." *Journal of Financial Economics*, 107(3): 715–739.

Mas, Alexandre. 2006. "Pay, reference points, and police performance." *The Quarterly Journal of Economics*, 121(3): 783–821.

Moffitt, Robert. 1990. "The econometrics of kinked budget constraints." *The Journal of Economic Perspectives*, 4(2): 119–139.

Odean, Terrance. 1998. "Are investors reluctant to realize their losses?" *The Journal of Finance*, 53(5): 1775–1798.

Saez, Emmanuel. 2010. "Do taxpayers bunch at kink points?" *American Economic Journal: Economic Policy*, 2(3): 180–212.

Samuelson, William, and Richard Zeckhauser. 1988. "Status quo bias in decision making." *Journal of Risk and Uncertainty*, 1(1): 7–59.

Sugden, Robert. 2003. "Reference-dependent subjective expected utility." *Journal of Economic Theory*, 111(2): 172–191.

Sydnor, Justin. 2010. "(Over) insuring modest risks." *American Economic Journal: Applied Economics*, 2(4): 177–199.